

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

ADAPTIVE POWER CONTROL

Inventors:

Sameer Halepete
H. Peter Anvin
Zongjian Chen
Godfrey P. D'Souza
Marc Fleischmann
Keith Klayman
Thomas Lawrence
Andrew Read

Prepared by:

Stephen L. King
30 Sweetbay Road
Rancho Palos Verdes, California 90275
(310) 377-5073

ADAPTIVE POWER CONTROL

BACKGROUND OF THE INVENTION

Field Of The Invention

This invention relates to computer systems and, more particularly, to
5 methods for varying the amount of power used by such systems during
use of the systems.

History Of The Prior Art

A significant problem faced by battery powered computers is the length
of time such computers are capable of operating between charges. As
10 computer processors become more capable, they tend to run at faster
speeds and dissipate more power. At the same time, the size and weight
of portable computers is constantly being reduced to make them more
portable. Since batteries tend to be a very significant element of the
weight of portable computers and other portable devices, the tendency
15 has been to maintain their size and thus their capacity at a minimum.

A typical portable computer today has an average life of approximately
two and one-half hours until its originally-full battery must be recharged.

A great deal of research has been directed to ways for extending the
operating life of portable computers. Presently, typical processors
20 include circuitry and software for disabling various power-draining
functions of portable computers when those functions are unused for
some extensive period. For example, various techniques have been
devised for turning off the screen when it has not been used for some

selected period. Similar processes measure the length of time between use of hard drives and disable rotation after some period. Another of these processes is adapted to put a central processor into a quiescent condition after some period of inactivity.

5 In general, these processes are useful in extending the operating life of a portable computer. However, the life still does not extend significantly beyond two and one-half hours for any computer having significant capabilities.

There has been a significant amount of research conducted from which 10 processor requiring less power might be produced. Most processors used in computer systems today are made using CMOS technology. The power consumed by a CMOS integrated circuit is given approximately by $P = CV^2f$, where C is the active switching capacitance, V is the supply voltage, and f is the frequency of operation. The maximum allowable 15 frequency is described by $f_{max} = kV$, where k is a constant.

It is desirable to operate the processor at the lowest possible voltage at a frequency that provides the computing power desired by the user at any given moment. For instance, if the processor is operating at 600 MHz, and the user suddenly runs a compute-intensive process half as 20 demanding, the frequency can be dropped by a factor of two. This means that the voltage can also be dropped by a factor of two. Therefore, power consumption is reduced by a factor of eight. Various methods of implementing this dynamic voltage-frequency scaling have been described in the prior art. All of these involve a component separate from 25 the processor on the system that provides multiple frequencies to

multiple system components. Also, they involve state-machines or power-management units on the system to coordinate the voltage-frequency changes. The efficiency of voltage frequency scaling is reduced when the frequency generator is not on the processor. Having a separate power-management unit increases the number of components in the system and the power dissipated by the system. It is also desirable to have the processor control both the voltage it receives and the frequency it receives. As the level of integration increases in processors, they control most of the system clocks; and it is desirable to provide control to the processor to change these clocks so they can be run at just the right frequency. Having a separate clock generator that produces multiple frequencies is not desirable because of the lack of tight coupling.

It is desirable to increase significantly the operating life of portable computers and similar devices.

15 Summary Of The Invention

It is, therefore, an object of the present invention to increase significantly the operating life of portable computers.

This and other objects of the present invention are realized by a method for controlling the power used by a computer including the steps of utilizing control software to measure the operating characteristics of a processor of the computer, determining when the operating characteristics of the central processor are significantly different than required by the operations being conducted, and changing the operating characteristics of the central processor to a level commensurate with the operations being conducted.

These and other objects and features of the invention will be better understood by reference to the detailed description which follows taken together with the drawings in which like elements are referred to by like designations throughout the several views.

5 Brief Description Of The Drawings

Figure 1 is a block diagram of various hardware components of a computer system utilized in accordance with the present invention.

Figure 2 is a flow chart illustrating the operation of one embodiment of the invention.

10 Figure 3 illustrates a number of registers utilized in the hardware components of the system shown in Figure 1.

Figure 4 is a block diagram illustrating the operation of sequencer circuitry which is a part of a processor illustrated in the system of Figure 1.

15 Detailed Description

Figure 1 is a block diagram of various hardware components of a computer system utilized in accordance with the present invention to control the operating frequency and voltage of the system. The hardware includes a processor 10, a clock generator 11, a programmable voltage generator 12, system memory (DRAM) 14, and an external battery (or other power supply) 13. The processor 10, clock generator 11, and voltage generator 12 are all mounted to a circuit board 15 in a manner known to those skilled in the art. The battery 13 and system memory 14

may be electrically connected to the circuit board in a number of possible ways known to those skilled in the art.

The processor 10 includes on the same semiconductor chip a number of components including a processing unit 16 and a programmable frequency generator 17. The processor 10 also typically includes a number of other components which are known to those skilled in the art but are not pertinent to the present invention and are therefore not illustrated. The processing unit 16 includes a number of logical components including a master control unit 18 which is the central portion for accomplishing clock and voltage control. In the present invention, the master control unit 18 also includes circuitry for monitoring the operating characteristics of the processor. Various monitoring functions (such as circuitry for accomplishing voltage and frequency monitoring) which are well known to the prior art are included as a part of the logical master control unit 18. The logical unit 18 may also include circuitry for making available additional information detected by other portions of the computer system in either analogue or digital form (e.g., temperature data). The logical unit 18 also includes circuitry for detecting other operations of the system including commands to be executed from which a particular type of operation to be executed may be determined. A detailed discussion of circuitry for providing various operating characteristics is included in U. S. patent application serial number 09/417,930, entitled Programmable Event Counter System, B. Coon et al, filed October 13, 1999, and assigned to the assignee of the present application.

The programmable frequency generator 17 receives an external frequency often referred to as a "slow clock" from the external clock generator 11.

The generator 17 responds to values furnished by control software executing on the processor to produce from the slow clock a core clock

5 for operation of the processing unit 16, one or more clocks for operation of the various system memory components shown as system memory 14 in the figure, the system bus, and any other components which might utilized a separate clock.

It should be specifically noted that contrasted to prior art systems, the

10 programmable frequency generator is able to provide individual

frequencies selectable for each of these components. Thus, prior art arrangements utilize an external clock generator to provide all of the different frequencies utilized by the system. This has a number of effects which are less than desirable. Since the clocks are generated off-chip, the time needed to change frequency is long. Since in an integrated processor all clocks are created from a single slow clock off chip, if the core frequency changes all of the frequencies change with it. Thus, a frequency furnished a single component cannot be changed without affecting a change in other frequencies. The voltage furnished by the

20 external clock generator does not change even though reduced

frequencies adapted to provide reduced levels of operations are furnished for various components of the system. A number of other factors slow the response of the system to changes in the various clocks when an external clock is used to generate the various operating frequencies for a system.

25

The core frequency for the processing unit 16 is generated by multiplying the slow clock by a factor. This factor is computed by the control software of the present invention which monitors the operation of the processor to determine from the characteristics of the processor just 5 what frequency should be selected. The manner in which the monitoring is accomplished and the effect it has on the control of the operating characteristics is described in detail below.

The frequencies at which the other components of the system operate are determined from the core frequency determined by multiplying the slow 10 clock by the core processor factor. For example, a system input/output (I/O) bus typically functions at a much slower frequency than does the processing unit. In the present invention, the control software computes the bus frequency by dividing the core frequency by a value. The process may also be conducted as a table lookup of an already computed value.

15 If the processing unit is conducting its current operations at a normal speed of 400 MHz, a bus frequency of 100 MHz. is derived by dividing the core clock by four. On the other hand, if the processing unit is capable of accomplishing its current operations at a relatively slow speed of 200 MHz, a bus frequency of 100 MHz. is still desirable since bus operations 20 are often the limiting factor in processing operations. In such a case, the control software computes a value of two as the divisor to obtain the bus frequency. It should be noted that although the bus frequency under discussion has been the system I/O bus, the invention may also be used for precisely choosing the operating frequencies for other system buses.

25 Similarly, various processors are often capable of utilizing system memory having different characteristics one of which is switching speed.

A system may utilize a plurality of interfaces between the processing unit and system memory in order to provide different operating frequencies for system memory which is being utilized. The present invention allows this to be easily accomplished by utilizing different divisors to obtain 5 different values from which the operating frequencies for different system memory units are determined. As will be noted in the following discussion, two different memory frequencies as utilized and more are possible.

Thus, by utilizing the phase-lock-loop generator 17 to determine a core 10 clock frequency and dividing that frequency by a plurality of different values determined by the control software, the operating frequencies for the different components of the system may be individually controlled and furnished to other components of the processor without the necessity of crossing chip boundaries with the consequent slowing 15 caused by negotiating the boundaries.

In order to allow the master control unit 18 to accomplish these operations, the processing unit 16 includes a number of registers which are utilized by the control software and the hardware. These include a master control register 20, a master status register 21, and a master 20 clock divider register 22 which are illustrated in Figure 3.

Of these registers, the clock divider register 22 stores, among other things, the multiplier computed by the control software for generating the core frequency, the value used as a divisor to obtain the bus frequency from the core frequency, a value used as a divisor to obtain a first system memory frequency from the core frequency, and a value used 25

as a divisor to obtain a second system memory frequency from the core frequency. In addition, the clock divider register 22 stores values used for various other including an indication that a frequency change command has been received.

5 The master control register includes values pertinent to the present description including the voltage which is to be furnished to the processor as a part of the change of frequency. This register also stores a value indicating the time period allowed for accomplishing the phase-lock-loop relock operation. The master status register also stores the 10 various values used as dividers and the value used as a multiplier to obtain the core frequency along with other significant information.

The various values stored in these registers are utilized, among other things, to control the operations of sequencer circuitry (illustrated in Figure 4) which carries out the operations necessary to changing the 15 frequency at which the components of the system operate. The sequencer circuitry carries out the series of steps required by which the phase-lock-loop circuitry is brought to the new frequency and relocked after the processor clock has been shut off.

The operations carried out by the sequencer commence at an idle state 20 which represents the normal condition of the sequencer in the absence of a frequency change operation. When the change frequency command and values are received, the sequencer steps from the idle condition to first shut down the core clock and the clocks to the various memory interfaces. The sequencer then waits a few cycles before shutting down 25 the bus clock, the master control clock, and saving information sufficient

to assure that timing during and after the sequencing is correct. After this delay, the sequencer starts a counter to time the phase-lock-loop relock process. When this count is complete, the sequencer wakes the bus and the master control units. Finally, the sequencer wakes up the core and memory interfaces and awaits another frequency changing 5 operation. The relation of the sequencer to the control software will be described in detail in the discussion of the process of the control software which follows.

Figure 2 is a flow chart representing the process carried out by one 10 embodiment of the invention. In the figure, the steps described in the left column represent operations accomplished by the control software, while the steps described in the right column represent operations accomplished by the cooperating hardware.

In a first step, the control software monitors various conditions of the 15 processor which relate to power expenditure by the processor. These conditions may include any of those described above including the present frequency and voltage of operation, the temperature of operation, the amount of time the processor spends in one of what may be a number of idle states in which various components of the system are 20 quiescent. For example, if the processor is running in what might be termed its normal mode of operation at a core frequency of 400 MHz. and a voltage of 1.3 volts, the control software may be monitoring the amount of time the processor spends in the "halt" state, the amount of time the processor spends in the "deep sleep" state, and the temperature of the 25 processor. The deep sleep state is a state in which power is furnished only to the processor and to DRAM memory. In this state, the processor

are all off and it does not respond to any interrupts. The halt state is a state in which the core clock has been stopped but the processor responds to most interrupts. If the processor is spending more than a preselected increment of its operation in these states while operating at 5 normal frequency and voltage, then power is being wasted. The detection of such operating characteristics therefore may indicate that the frequency and voltage of operation should be reduced.

On the other hand, it may be found that the processor is functioning at a reduced frequency and voltage and that a series commands have been 10 furnished to be executed by the processor which require greater processing power. In such a case, these characteristics suggest that it may be desirable to increase the voltage and frequency of operation in order to handle these commands.

Consequently, the control software detects operating characteristics and 15 determines whether those characteristics indicate that the frequency and voltage of operation should be changed. From the possible sets of conditions, the control software detects the particular set involved and computes correct values for the core clock frequency, the core clock frequency multiplier, the various DRAM clock frequency dividers, and the 20 bus frequency divider. If any other components of the circuitry receive their own clocks, then multipliers or dividers for these values are computed. It should be noted that the control software may actually compute the various values required for the given characteristics which have been determined or may utilize a lookup table storing precomputed 25 values.

At a next step, the software reviews the values computed and determines whether the frequency is to be increased. If the frequency is to be increased, it is first necessary that the voltage be increased to allow the processor to function at a higher frequency. In such a case, it is first

5 necessary to increase the voltage level of operation. The typical power supplies offer a number of pins (often five) by which different operating voltages may be selected. This allows a range of different voltages to be provided. Consequently, the control software simply furnishes a correct value on the input pins of the power supply to cause the computed voltage to be furnished to the frequency generator and to the processor.

10 In one embodiment, the voltage increase is accomplished by providing a level to be reached and a time period for the voltage to settle to this level.

15 It should be noted that the voltage may be increased in a single step, an action which would typically cause phase-locked-loop circuitry of a frequency generator to lose its lock and would create a large surge of current causing the currently-available voltage regulator circuitry to initiate a system reset. This problem may be eliminated with future voltage regulator circuitry. Alternatively, the voltage may be increased in a series of small steps which would not have this effect. For example, if

20 increases of approximately 50 millivolts are enabled, then the frequency generator will remain stable during the voltage increase and a system reset will not occur. This offers the advantage that the processor may continue to execute commands during the period in which the voltage change is taking place.

25 If the control software was not increasing but rather decreasing frequency of operation at the previous step, then the original voltage level

is not changed at this time. In either case, the control software then goes through a sequence of steps in which various operations of the processor are prepared for shutdown so that the system clocks can be changed. With a particular processor such as that referred to in the 5 patent application described above, this includes flushing a gated store buffer, suspending bus and direct memory access (DMA) operations, and enabling self-refreshing circuitry for system DRAM memory.

With these processor operations shut down, the control software transfers the new divider values and writes a bit indicating a frequency 10 change is to occur. The hardware stores the divider values in the clock divider register and the change frequency indicator. This starts the hardware process of the sequencer. The control software then writes “stop core,” “stop DRAM0” and “stop DRAM1” bits of the master command register to stop the clocks being furnished to these 15 components.

Writing the master control register bits to stop the clock frequencies and the values to the hardware causes the hardware to commence the remainder of the frequency changing operation utilizing the sequencer circuitry described above. At this point, the software effectively goes into 20 a wait state which continues until the core clock is enabled at the new frequency. The sequencer responds to the command by shutting down the core clock and the DRAM memory interfaces. The sequencer pauses for sufficient time to assure that this has happened and then shuts down the bus and master control clocks.

Because the core clock has been stopped, timing must be accomplished based on the external clock furnished to the system during this period. Counter circuitry dependent only on phase-lock-loop relock time is utilized to measure the time allowed for the phase-lock-loop circuitry to lock to the new frequency. At this point, the sequencer utilizes the new values furnished to effect a new value for the core (and other) frequency. After a safe lock period has passed ("relock time" stored in the master control register), the sequencer wakes the bus and master control units. The sequencer waits a few clocks of the slow frequency and then turns on the core clock and the DRAM interfaces.

Because the internal clocks of the system are shut down during the operation of the sequencer, it is necessary that the system provide a means of maintaining timing consistent with the normal world clock. Computer systems utilize a time stamp counter to keep track of world clock values. The value kept in this counter is utilized for certain operations conducted by the central processing unit. Once the phase-lock-loop circuitry of the frequency generator 17 has been stopped, the value in the time clock counter no longer represents accurate world time. Moreover, when the new frequency is reached and locks in, the rate at which the counter is iterated will change. To provide for accurate time stamp readings, a number of lower-valued bits indicating the last time of program execution held by the time stamp counter are stored. These are furnished to the control software along with the relock time value and the new frequency once the frequencies have restabilized to allow accurate computation of the normal world time.

25

Once the clocks have been turned on at the new frequencies, the control software ends its wait state and determines whether the operation was to decrease the frequency. Assuming the operation was to increase the frequency, the software then recalculates the time stamp counter value 5 and checks the various interface timings to assure that they are correct. If the operation was to decrease the frequency, the control software causes the voltage to be lowered to the calculated value (either in one or a series of incremental steps) and then recalculates the value for the time stamp counter and checks the interface timings. At that point, the 10 control software begins again to monitor the various conditions controlling the frequency and voltage of operation.

It should be noted that at some point during the monitoring operation it may be found that the processor is functioning at a normal frequency and voltage, that the temperature of operation is below some preselected 15 value, and that a series of processor-intensive commands have been furnished to be executed by the processor. In such a case, these characteristics suggest that it may be desirable to increase the voltage and frequency of operation in order to handle these commands for a period less than would raise operating temperatures beyond a safe level. 20 In such a case, the control software may compute higher frequency and voltage values and a temperature (or a time within which temperature will not increase beyond a selected level) in order to cause the hardware to move to this higher frequency state of operation. In such a case, the processor executing the process illustrated effectively ramps up the 25 frequency and voltage so that the processor "sprints" for a short time to accomplish the desired operations. This has the effect of allowing a

processor which nominally runs at a lower frequency to attain operational rates reached by more powerful processors during those times when such rates are advantageous.

5 Although the present invention has been described in terms of a preferred embodiment, it will be appreciated that various modifications and alterations might be made by those skilled in the art without departing from the spirit and scope of the invention. The invention should therefore be measured in terms of the claims which follow.

What Is Claimed Is: